

Initial Congestion Window Specification

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on May 15, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the BSD License.

Abstract

This document specifies the progression of initial TCP congestion window sizes over the next nine years.

1 Introduction

This document proposes a series of increases to TCP's [\[RFC790\]](#) initial congestion window. For the first time in roughly nine years the IETF is currently considering increasing the initial congestion

Expires: May 15, 2011

[Page 1]

window (IW). The current proposal is to increase IW from the standard two to four segments [[RFC5681](#)] to ten segments [[CDCM10](#)]. While there is much evidence ([[Chu09](#),[CDCM10](#)] and references therein) that the increase has benefit, there is also some unease within the community that stems from a lack of solid understanding of the dynamics that the increase would cause system-wide. Such an understanding is nearly impossible to apprehend. This document presents an alternate approach that slowly increases the size of the allowed IW over the next nine years---attaining a value of ten segments over the course of the first three years and increasing to 15 segments by the end of the inflationary period. By making slow and steady changes the community can continually assess the dynamics and short-circuit the increases as necessary. We believe this offers the best tradeoff between (1) assuring network safety and (2) allowing for reasonable performance increases as network capacity increases.

2 Initial Congestion Window Values

This document specifies experimental and standard values for the initial window as a function of time in the following chart.

Year	Exp. IW	Std. IW
pre-1998		(1, 1*SMSS)
1998	(2--4, 4380) [RFC2414]	(1, 1*SMSS)
1999		(2, 2*SMSS) [RFC2581]
2002		(2--4, 4380) [RFC3390]
2009		(2--4, 4380) [RFC5681]
2011	(6, 8760)	
2012	(8, 11680)	(6, 8760)
2013	(10, 14600)	(8, 11680)
2014	(11, 16060)	(10, 14600)
2015	(12, 17520)	(11, 16060)
2016	(13, 18980)	(12, 17520)
2017	(14, 20440)	(13, 18980)
2018	(15, 21900)	(14, 20440)
2019		(15, 21900)

The (X,Y) tuples indicate the IW in terms of the maximum number of segments, X, and the maximum number of octets, Y. The smaller of these two values is used as the IW. In other words, (6,8760) indicates that six segments of 1460 bytes each can be sent. If the packet size is larger than 1460 bytes then the IW is bounded at 8760 bytes. If the segment size is smaller than 1460 bytes then the IW is bounded by six segments.

Each IW value spends one year as experimental and assuming the community does not find any broad problems because standard the

following year.

We take larger steps at the beginning of the time period than at the end due to our confidence in each step. That is, [[CDCM10](#)] (and references therein) shows that an IW of ten segments is reasonably

Expires: May 15, 2011

[Page 2]

safe today. Therefore, our first step to six segments seems conservative, as does the overall progression to ten segments over the next three years. After that the empirical basis for further increases, which exists in some form today [[CDCM10](#)], is less compelling and therefore further steps are increases of only a single segment at a time.

We stress that the above is an upper bound on the allowable IW and not a requirement to use the given value. Hosts may indeed have reasons to be less aggressive in certain situations.

3 Short Circuiting

The table in the last section can be short-circuited if the IETF finds particular issues with a given IW. While our expectation is the IW values given in [Section 2](#) are and will be safe for general Internet use the progression can be canceled by making this document obsolete.

4 Security Considerations

This document discusses the initial congestion window permitted for TCP connections. Changing this value does not raise any known new security issues with TCP.

5 IANA Considerations

None

Normative References

[RFC793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC2414] Allman, M., Floyd, S. and C. Partridge, "Increasing TCP's Initial Window Size", [RFC 2414](#), September 1998.

[RFC2581] Allman, M., Paxson V. and W. Stevens, "TCP Congestion Control", [RFC 2581](#), April 1999.

[RFC3390] Allman, M., Floyd, S., C. Partridge, "Increasing TCP's Initial Window", [RFC 3390](#), October 2002.

[RFC5681] Allman, M., Paxson V. and E. Blanton, "TCP Congestion Control", [RFC 5681](#), September 2009.

Non-Normative References

[Chu09] Chu, J., "Tuning TCP Parameters for the 21st Century",
<http://www.ietf.org/proceedings/75/slides/tcpm-1.pdf>, July 2009.

Expires: May 15, 2011

[Page 3]

[CDCM10] Chu, J., Dukkupati, N., Cheng, Y. and M. Mathis,
"Increasing TCP's Initial Window", Internet-Draft
[draft-ietf-tcpm-initcwnd-00.txt](#) (work in progress), October
2010.

Author's Addresses

Mark Allman
ICSI
1947 Center Street
Suite 600
Berkeley, CA 94704-1198

Phone: 440-235-1792
EMail: mallman@icir.org
<http://www.icir.org/mallman/>

Expires: May 15, 2011

[Page 4]